The 203 patent (Patent No. 6,484,203) – GrIDS invalidates

The claims asserted against ISS from the 203 patent are 1-6 and 12-17. My opinion is that GrIDS falls entirely within the patent claims of the 203 patent, showing all elements of the claims, with the exception of two dependent claims (which are materially identical to each other). Those I believe are invalid for other reasons. Thus I believe the patent claims asserted in suit are entirely invalid. I attach charts at Tab A showing that the GrIDS references disclose the elements of the claims.

I will take the claims in sequence, and argue in the alternative based on the three proposed claim constructions where necessary.

> *1. A computer-automated method of hierarchical event monitoring and analysis within an enterprise network comprising:*
>
> *deploying a plurality of network monitors in the enterprise network;*
>
> *detecting, by the network monitors, suspicious network activity based on analysis of network traffic data selected from the following categories: {network packet data transfer commands, network packet data transfer errors, network packet data volume, network connection requests, network connection denials, error codes included in a network packet};*
>
> *generating, by the monitors, reports of said suspicious activity; and*
>
> *automatically receiving and integrating the reports of suspicious activity, by one or more hierarchical monitors*

The Grids graph engines as described above were deployed in enterprise networks, and they analyzed data about the network traffic in at least two of the categories above: network connection requests and network connection denials (Tab A). They detected suspicious network activity by doing so, and they generated reports of that activity (in the form of GrIDS graphs as described above). Those reports were automatically received and integrated by one or more hierarchical graph engines.

It seems clear to me that under all the proposed constructions, shown in the Joint Claim Construction Chart, GrIDS met all elements of this independent claim.

SRI construes "monitor" broadly to be "any process or component that can analyze data", either network traffic data, or reports of suspicious activity. GrIDS graph engines analyzed data of both of those kinds so clearly are monitors in that sense.

However, SRI's construction of the terms "monitor" and "hierarchical monitor" is inconsistent with the definition of monitor in the specification defines monitor as follows:

> *All monitors (service, domain, and enterprise) 16a-16f use the same monitor code-base. However, monitors may include different resource objects 32 having different configuration data and methods. Thus reusable software architecture can reduce implementation and maintenance efforts. Customizing and dynamically configuring a monitor 16 thus becomes a question of building and/or modifying the resource object 32.* 1203, 29-37.

SRI defines as monitor "Process or component in a network that receives reports from at least one lower-level monitor". These constructions imply that a hierarchy may only have two levels, and the monitors at each level can be anything that analyzes either network traffic data or reports of suspicious activity. That is inconsistent with the monitors defined in the specification.

For ISS, a "monitor" (or interchangeably "network monitor") is "generic code that can be dynamically configured and reconfigured with reusable modules that define the monitors inputs, analysis engines, and their configurations, response policies and output distribution for its reports." [Joint Claim Construction Statement]. This is the monitor defined in the specification.

Monitors under ISS's construction are present in GrIDS. A GrIDS graph engine meets all the tests for being a monitor in the ISS sense. The exact same code was used at every level in the hierarchy. The hierarchy of those engines could be "configured and reconfigured", which would define the "inputs" and "outputs" of the monitors. GrIDS rulesets essentially function like independent reusable analysis engines that could be inserted or deleted from the system, and controlled the configuration, response policies and output of reports.

Symantec's definition of "monitor" (or interchangeably "network monitor") is "software that can be dynamically configured to collect, analyze and respond to suspicious network activity, and that included one or more analysis engines and a resolver that implements a response policy." GrIDS was software that could be dynamically configured to collect, analyze and respond to suspicious network activity. It included an analysis engine. In addition, GrIDS did allow a user to specify a variety of possible responses, and thus a section of the ruleset code specified what to do if a graph crossed some threshold for being suspicious. That ruleset code could be considered a resolver that implements a response policy.

In summary, I cannot see how independent claim 1 of the 203 patent can possibly be regarded as novel over the prior art, under any of the three constructions.

Next we consider the dependent claims that refer back to claim 1.

2. The method of claim 1, wherein integrating comprises correlating intrusion reports reflecting underlying commonalities.

GrIDS clearly did this. The graph engines combined graph reports based on underlying commonalities such as the fact that the two graphs shared a common host, the two graphs were close together in time, and so forth. These commonalities are "relationships among the reports". Hence if claim 2) is combined with claim 1), GrIDS still preempts the 203 patent (under all proposed constructions of the terms).

3. The method of claim 1, wherein integrating further comprises invoking countermeasures to a suspected attack.

To assess this, we clearly need to know what exactly are "countermeasures". SRI and Symantec are very broad: "Taking an action in response". ISS is equally broad but spells out some possibilities: "taking an action in response to a suspected attack, including passive responses such as report dissemination to other monitors or administrators, and highly aggressive actions, such as severing a communication channel of the reconfiguration of logging facilities within network components."

There is support for the ISS version within the specification of the patent itself. In column 11, lines 32-43, we find:

> "Countermeasures range from very passive responses, such as report dissemination to other monitors 16a-16f or administrators, to highly aggressive actions, such as severing a communication channel or the reconfiguration of logging facilities within network components (e.g., routers, firewalls, network services, audit daemons). An active response may invoke handlers that validate the integrity of network services or other assets to ensure that privileged network services have not been subverted. Monitors 16a-16f may invoke probes in an attempt to gather as much counterintelligence about the source of suspicious traffic by using features such as traceroute or finger."

In GrIDS, as discussed earlier, responses were taken in the assessment rules of a ruleset. [GRIDS section 2.3.3] Built-in functions for taking responses included the *alert* and *response* functions; these automatically sent reports to the user interface for administrator review. Additionally, the ruleset controls whether reports are sent to higher level monitors or not [GrIDS section 2.4.3]. Thus GrIDS could clearly "invoke countermeasures to a suspected attack" of a passive kind.

Further, since assessment rules can call user defined functions [GRIDS section 2.3.6], the response could be anything a user can code in the Perl language, which allows for basically anything a computer can do. Aggressive responses to cut communication and kill processes were in wide discussion in the intrusion detection community prior to November 1997. For example it was explicitly disclosed by SRI in detail in the 1997 NISSC paper [EMERALD p.361], and was also extensively presented to the DARPA PI community in presentations on the IDIP protocol [IDIP]. Thus GrIDS explicitly was capable of taking passive responses, which appear to be included in the definition of "countermeasures" and it would have been obvious to one skilled in the art that GrIDS could be extended to take aggressive responses.

Therefore, if we combine dependent claim 3 with independent claim 1 of the 203 patent, GrIDS falls entirely within it, and thus these claims are not novel over prior art.

> *4. The method of claim 1, wherein the plurality of network monitors include an API for encapsulation of monitor functions and integration of third-party tools.*

GrIDS includes an API for incorporating arbitrary user defined functions into the graph engine [GrIDS section 2.3.6]. This allows arbitrary Perl functions and thus would easily support either taking actions or acquiring additional data from other programs or systems,

or allowing other systems to take in GrIDS data for further processing. Furthermore, GrIDS by design was intended to incorporate information from a range of data-sources, and to provide a standard protocol for which numerous different types of source data could be expressed. As the *Design of GrIDS* report puts it in section 1.1.2:

> "Node and edge attributes may come from other IDSs, network sniffers, or any monitor that is equipped with a filter to send its output to GrIDS. A well defined syntax for reporting to GrIDS will be available to GrIDS users who wish to write their own filters."

Thus GrIDS is a system which would fall under claim 1 as refined by claim 4 of the 203 patent. Thus this is not novel over prior art either.

Additionally, in my opinion, adding APIs to export and import functionality of software systems is a widespread and standard practice in software engineering and should have been obvious to anyone with any professional knowledge of software development.

Furthermore, SRI was participating in the CIDF process, an open public standards effort tasked to provide APIs for intrusion detection modules to allow them to interoperate, and which began this task before the statutory bar date for these patents, which was well known in the field at the time. SRI should not be simultaneously participating in a public process to develop APIs for intrusion detection modules, and then going back after the fact and attempting to patent the idea of adding APIs to intrusion detection modules.

5. The method of claim 1, where the network is a TCP/IP network.

The implementation of GrIDS took data from TCP/IP networks – indeed that was the only kind of network it took data from. GrIDS was intended to be deployed on TCP/IP enterprise networks. Thus GrIDS is a system that would fall under claim 1 as refined by claim 5 of the 203 patent. Thus this is not novel over prior art.

6. The method of claim 1, wherein the network monitors are deployed at one or more of the following facilities of the enterprise network: {gateways, routers, proxy servers}.

GrIDS was not applied directly to data from gateways, routers, and proxy servers. However, GrIDS by design was intended to incorporate information from a range of systems, and to provide a standard protocol for which numerous different types of source data could be expressed. As the *Design of GrIDS* report puts it in section 1.1.2:

> "Node and edge attributes may come from other IDSs, network sniffers, or any monitor that is equipped with a filter to send its output to GrIDS. A well defined syntax for reporting to GrIDS will be available to GrIDS users who wish to write their own filters."

The deployment of intrusion detection algorithms in at least routers was under broad public consideration in the intrusion detection community prior to November 1997,

including and especially at DARPA principal investigator meetings where SRI staff were present. For example, the IDIP system [IDIP] had modules specifically designed to operate in routers and have the routers collaborate to trace the source of problems and block connections from systems that had been identified as intrusive. The JiNao system had modules intended to work inside routers [JiNao]. Chen and Levitt described a protocol for doing intrusion detection in routers to detect misbehaving ones. That paper was published in September 1997 [Che97]. SRI themselves disclosed the possibility prior to the statutory bar. For example, in the 1997 NISSC paper they write on page 355:

> "Service monitors are dynamically deployed within a domain to provide localized real-time analysis of infrastructure (e.g., routers or gateways)…"

Thus it would have been obvious to one skilled in the art in 1997 that GrIDS could incorporate data from routers, and it would have been a basically trivial matter to write such a filter to convert router logs or packet data to the GrIDS message format. With these obvious extensions, GrIDS is a system that would fall under claim 1 as refined by claim 6 of the 203 patent. Thus this is not novel over prior art.

Finally, claims 12 to 17 are materially identical to claims 1 to 6 except that they are "system" claims rather than "method" claims. It appears to me that this has no influence on the fact that GrIDS fits within all of them. For the sake of brevity and clarity, I will not repeat the arguments made above for claims 1 to 6, but the reader may apply them to claims 12-22 by simply incrementing the numbers above by 11 if necessary.

The 615 patent (Patent No. 6,711,615)

The 615 patent claims asserted against ISS (1-6 and 13-18) do not differ materially from the 203 patent claims asserted. After a short discussion of the differences in Claim 1, I simply cross-reference back to the relevant analysis of the 203 claims, rather than repeat an identical discussion.

> *1. A computer-automated method of hierarchical event monitoring and analysis within an enterprise network comprising:*
>
> *deploying a plurality of network monitors in the enterprise network;*
>
> *detecting, by the network monitors, suspicious network activity based on analysis of network traffic data selected from one or more of the following categories: {network packet data transfer commands, network packet data transfer errors, network packet data volume, network connection requests, network connection denials, error codes included in a network packet, network connection acknowledgements, and network packets indicative of well-known network-service protocols};*
>
> *generating, by the monitors, reports of said suspicious activity; and*
>
> *automatically receiving and integrating the reports of suspicious activity, by one or more hierarchical monitors.*

The only material difference from claim 1 of the 203 patent is that the list of possible categories of network traffic has been made longer by the addition of "network connection acknowledgements", and "network packets indicative of well-known network-service protocols". Since GrIDS fit entirely in the shorter list, a-fortiori it fits in the longer list. Indeed of the added categories GrIDS definitely monitored at least "network packets indicative of well-known network-service protocols". For example, there is a detailed discussion in the GrIDS design report of monitoring telnet (section 7.3.1) and NFS (section 7.3.2) and turning it into GrIDS graph reports (the GrIDS equivalent of "events" in Emerald or DIDS). Some quotations from the telnet section should suffice to establish the level of detail:

> "A TELNET connection has several stages (as shown in Figure 7.1). The stages are START, OPTION_NEGOTIATION, AUTHENTICATION, DATA, END or RESET. Each stage is reported by the sniffer with a connection report. We describe the attributes included with each report.
>
> (TELNET,TCP,END,SUCC) refers to the event that FIN packets were observed in both directions. (TELNET, TCP, END, SUCC) indicates that a TCP RST flag was sent in one direction. The other side of the TCP connection may continue to send data and delay acknowledgement.
>
> "In any case, the sniffer reports the first RST sent in either direction without waiting for its corresponding ACK message.

And a little later:

> "The reports are formatted in the DOT-like graph language (see Chapter 3). Here is a report of a START event:
>
> "Digraph sniffer {helvellyn.cs.ucdavis.edu->jaya.cs.ucdavis.edu [app-prot="telnet", prot="tcp", sport="1024", dport="23", stime=33311222, seq-12345, stage="start", status="succ"];}

This is followed by syntax details for END events and AUTH events for the telnet protocol. This establishes clearly that GrIDS was basing its analysis, in part, on "network packets indicative of well-known network-service protocols" (i.e. telnet).

> Thus the 615 claim is invalid.

> *2. The method of claim 1, wherein integrating comprises correlating intrusion reports reflecting underlying commonalities.*

See the analysis of the 203 patent claim 2.

> *3. The method of claim 1, wherein integrating further comprises invoking countermeasures to a suspected attack.*

See the analysis of the 203 patent claim 3.

4. The method of claim 1, wherein the plurality of network monitors include an API for encapsulation of monitor functions and integration of third-party tools.

See the analysis of the 203 patent claim 4.

5. The method of claim 1, wherein the enterprise network is a TCP/IP network.

See the analysis of the 203 patent claim 5.

6. The method of claim 1, wherein the network monitors are deployed at one or more of the following facilities of the enterprise network: {gateways, routers, proxy servers}.

See the analysis of the 203 patent claim 6.

13. An enterprise network monitoring system comprising:
a plurality of network monitors deployed within an enterprise network, said plurality of network monitors detecting suspicious network activity based on analysis of network traffic data selected from one or more of the following categories: {network packet data transfer commands, network packet data transfer errors, network packet data volume, network connection requests, network connection denials, error codes included in a network packet, network connection acknowledgements, and network packets indicative of well-known network-service protocols};
said network monitors generating reports of said suspicious activity; and
one or more hierarchical monitors in the enterprise network, the hierarchical monitors adapted to automatically receive and integrate the reports of suspicious activity.

This is not materially different than claim 1 of this 615 patent. See the analysis there.

14. The system of claim 13, wherein the integration comprises correlating intrusion reports reflecting underlying commonalities.

See the analysis of the 203 patent claim 2.

15. The system of claim 13, wherein the integration further comprises invoking countermeasures to a suspected attack.

See the analysis of the 203 patent claim 3.

16. The system of claim 13, wherein the plurality of network monitors include an application programming interface (API) for encapsulation of monitor functions and integration of third-party tools.

See the analysis of the 203 patent claim 4.

17. The system of claim 13, wherein the enterprise network is a TCP/IP network.

See the analysis of the 203 patent claim 5.

18. The system of claim 13, wherein the network monitors are deployed at one or more of the following facilities of the enterprise network: {gateways, routers, proxy servers}.

See the analysis of the 203 patent claim 6.

This completes my analysis of the asserted claims of the 615 patent. All are invalid.

Response to SRI's contention on GrIDS

In their response to ISS's second set of interrogatories, SRI spends one page arguing that GrIDS does not invalidate the asserted claims of the 203 and 615 patents.

They make three independent points, which we shall refute in turn.

Firstly they assert that it is unproven that the "GrIDS 1997" reference [GrIDS] was publicly available prior to the statutory bar date. As discussed above, I have specific memory of having made it available on the web by the time of the August 1996 Santa Cruz PI meeting. Additionally, the Wayback Machine recorded a copy of the U.C. Davis GRIDS webpage link to the design report on July 17, 1997. That was the link I created to the design report. When I clicked on that the link during the preparation of the report it leads to the May 14, 1997 paper. Therefore, I conclude the "GrIDS 1997" was on the web prior to November 9th, 1997.

Secondly, SRI asserts that GrIDS graph engines cannot be a hierarchical monitor within the context of patent 203 claims 1 and 12, and patent 615 claims 1 and 13 because:

> "the higher level engines in the GrIDS system do not perform the required tasks of the hierarchical monitors recited in independent claims 1 and 12 of the '203 patent and claims 1 and 13 of the '615 patent. The only time information is sent to a higher level engine occurs when a Graph Engine sends a subset of the information in a graph to a higher level engine, and this information is only sent when a connection is made to a computer outside its protion of the network ("GrIDS 1997" 16). Therefore, the citations by the Defendants show that the higher level engines do not automatically receive and integrate reports of suspicious activity as recited in claims 1 and 12 of the '203 patent and claims 1 and 13 of the '615 patent. The Graph Engines at the lower levels of the hierarchy use their Rulesets to analyze the graphs they have built. They then create alerts and reports based on this analysis. ("GrIDS 1997" 20-22). However, rather than sending these reports to the higher level monitors, they are sent to that Graph Engine's log file. ("GrIDS 1997") 22."

This is factually incorrect. In fact, GrIDS always sent graphs that a lower graph engine in the hierarchy would consider suspicious to the graph engines above. This was done in order to ensure that higher level engines were always aware of identified suspicious activity, in addition to performing correlation on activities of uncertain suspicion level in order to detect distributed attacks.

This was done through the mechanism of global attributes of graphs, which stored variables that covered global properties of the graph such as the number of nodes, number of edges, etc. When global attributes of a graph changed, the graph would be propagated upwards. Thus any change in the size of a graph, for example, even in the absence of a connection to another department, would cause a node report for that graph in that department to be sent up to the parent department for incorporation into its graph building operations.

This operation is disclosed in the "GrIDS 1997" report. Specifically on p 17, it begins:

> "Recall that a department shows up as a node in the parent domain's aggregator and that rulesets may permit multiple instances of nodes with the same name to appear in separate graphs within a ruleset. Not only are graph attribute reports destined for a particular ruleset in the parent, they are also, in general, destined for a particular instance of a node.
>
> "Associated with each graph being constructed by the engine is a graph identifier (gid) that is unique within the ruleset for a department. Similarly, associated with each node that represents a department is an instance identifier. The instance id on a node is the same as the gid of the graph whose global attributes correspond to the node's attributes. Thus there is a correspondence between a graph at a lower level and a node at a higher level. When the global attributes on a graph change, the updated attribute values are sent to the particular node in the parent that correspond to the graph at lower department level (i.e. has the same name as that of the department and has the instance id that is that same as the gid). The attributes that are sent up are constructed into a graph (as per the rules in the ruleset) and merge (also as per ruleset rules) with the graph containing that instance, thereby (potentially) updating the attributes on a node and graph."

On p 13-14, the text refers to the auto-computed global attributes of a graph (ie those that were invariably calculated by the engine itself, not by the specific ruleset):

> "The following attributes are computed automatically by the engine. They may be referred to by the rule writer, but are not computed explicitly by the rules.
>
> - "Global Attributes
>   - "*gids*, a set of graph ids associated with this graph, any of which can be used as a unique identified
>   - "*ruleset*, the name of the ruleset this graph is in.
>   - "*nnodes*, the total number of nodes in a graph.
>   - "*nedges*, the total number of edges in a graph"

This makes it clear that since the size of the graph (in terms of either nodes or edges) was always a global variable, any change in it would invariably result in upward propagation of a reduced graph recording the change.

Additionally, in the example ruleset on page 11-13, there is a global attribute "alerts" explicitly computed by that ruleset which keeps track of whether any reports with the "alert" attribute set have been incorporated into the graph (and the list of textual values of those alerts). A change in this global attribute would trigger propagation of a node report upwards even in the absence of any connections at all. Since the node rules for making a report into the graph in this example would incorporate the alert attribute from any incoming node report with the "alert" attribute into the global alerts variable, this would cause a single report of an intrusion from a host monitor to propagate all the way to the top of the sub-hierarchy over which this particular ruleset was defined (which might be the entire hierarchy).

Thus SRI's contention "that the higher level engines do not automatically receive and integrate reports of suspicious activity" is without merit. GrIDS had a precise and configurable mechanism for determining under what circumstances graphs were suspicious enough to pass upwards and thus be automatically received and integrated into reduced graphs at the ancestor monitors. This was regularly done and is clearly disclosed in the "GrIDS 1997" reference as shown above.

The 203 patent (Patent No. 6,484,203) – DIDS invalidates

The claims asserted against ISS from the 203 patent are 1-6 and 12-17. My opinion is the claims analyzed below are invalid over DIDS under the claim construction of SRI, and also that of Symantec. In the following all references to DIDS refer to [DIDS] unless otherwise specified.

I will take the claims in sequence.

> *1. A computer-automated method of hierarchical event monitoring and analysis within an enterprise network comprising:*
>
> *deploying a plurality of network monitors in the enterprise network;*
>
> *detecting, by the network monitors, suspicious network activity based on analysis of network traffic data selected from the following categories: {network packet data transfer commands, network packet data transfer errors, network packet data volume, network connection requests, network connection denials, error codes included in a network packet};*
>
> *generating, by the monitors, reports of said suspicious activity; and*
>
> *automatically receiving and integrating the reports of suspicious activity, by one or more hierarchical monitors*

SRI construes "monitor" broadly to be any process or component that can analyze data, either network traffic data, or reports of suspicious activity.

With these interpretations of the terms, it is clear that DIDS has all of the elements of a system covered by claim 1. DIDS had a LAN monitor (the NSM) which could "analyze

network traffic data", and a plurality of those could be deployed in the enterprise network. On page 168 of [DIDS] it says that, "DIDS components included the DIDS director, a single host monitor per host, and a single LAN monitor for **each** LAN segment of the monitored network." (Emphasis added). This makes clear that a plurality of monitors was specifically contemplated in DIDS.

The DIDS paper also clearly discloses at least the analysis of "network packet data volume", and "network connection requests", for example on page 171 where the LAN monitor is disclosed to audit "host-to-host connections, services used, and volume of traffic per connection".

Clearly, the LAN monitor was intended to "detect suspicious network activity based on analysis of network traffic data": that was its entire purpose.

We also know that the LAN monitors in DIDS generated DIDS events, which constitutes "generating, by the monitors, reports of said suspicious activity" and utilized the LAN Agent subcomponent to send these to the DIDS director. The DIDS director is a monitor according to the SRI claim construction since it is a "process or component that can analyze data", and specifically it analyzes "reports of suspicious network activity" – those sent to it by the LAN Monitors (and also Host Monitors). Further, the DIDS director is a hierarchical monitor, since it "receives reports from at least one lower-level monitor", to whit the LAN Monitors.

Thus DIDS is engaged in "automatically receiving and integrating the reports of suspicious activity, by one or more hierarchical monitors" since the Director was certainly integrating the reports as it ran its rules on them, decided which activities were performed by the same user as which other activities on other computers, and reported suspected intrusions to the users of the system. Hence all elements of claim 1 of the 203 patent were present in DIDS publications seven years prior to the filing of the earliest of the patents in suit. Thus, under the SRI claim construction, this claim impermissibly covers prior art and is invalid.

Also, under Symantec's claim construction, a "monitor" (or interchangeably "network monitor") is "software that can be dynamically configured to collect, analyze and respond to suspicious network activity, and that included one or more analysis engines and a resolver that implements a response policy."

I believe that under this definition, both a DIDS director and a LAN monitor are "monitors" within the meaning of the patent. They were software that could be "dynamically configured". In the case of the LAN monitor this was because (p169)

> "The director can also make requests for more detailed information from the distributed monitors via a "GET" directive, and issue commands to have the distributed monitors modify their monitoring capabilities via a "SET" directive."

The DIDS director could also be "dynamically configured" because (p 169) we learn that the user interface would allow the security officer to configure the system by "setting "wire-taps", and requesting "more specific types of information from the monitors".

Both components were built specifically to "collect, analyze and respond to suspicious network activity" (given that responding can comprise simply reporting to another monitor or a system administrator). The LAN monitor had an "analysis engine" (the

LAN event generator), and a "resolver that implements a response policy" (the LAN agent that sends notifications to the director). Similarly, the Director had an analysis engine in the form of its expert system and code to decide whether or not to report the activity, which constituted the resolver.

The rest of the analysis is the same under either Symantec's or SRI's claim construction. Either way, claim 1 is invalid.

Next we consider the dependent claims that refer back to claim 1.

2. The method of claim 1, wherein integrating comprises correlating intrusion reports reflecting underlying commonalities.

The patent specification itself discloses almost nothing about what is meant by "correlating intrusion reports reflecting underlying commonalities". The only relevant text I am able to find is in column 4 line 40:

> "The enterprise monitor 16f (or monitors, as it would be important to avoid centralizing any analysis) focuses on network-wide threats such as Internet worm-like attacks, attacks repeated against common network services across domains, or coordinated attacks from multiple domains against a single domain. As an enterprise monitor 16f recognized commonalities in intrusion reports across domains (eg. The spreading of a worm or a mail system attack repeated throughout the enterprise 10), the monitor.." (discussion on possible responses follows)

It does not appear to me that this text enables anyone to build a system capable of "correlating intrusion reports reflecting underlying commonalities". It announces that such functionality would be desirable, but gives no specifics on how it might be accomplished.

So, given that this is one of SRI's techniques of "correlating intrusion reports reflecting underlying commonalities", it seems to me clear that DIDS would have to fall under the same claim, on page 168 of [DIDS] we find

> "In a doorknob attack the intruder's goal is to discover, and gain access to, insufficiently-protected hosts on a system. The intruder generally tries a few common account and password combinations on each of a number of computers."

And

> "Because DIDS aggregates and correlates data from multiple hosts and the network, it is in a position to recognize the doorknob attack by detecting the pattern of repeated failed logins even though there may be too few on a single host to alert that host's monitor."

Thus DIDS invalidates claim 2 under SRI's and Symantec's construction.

3. The method of claim 1, wherein integrating further comprises invoking countermeasures to a suspected attack.

To assess this, we clearly need to know what exactly are "countermeasures". SRI and Symantec are very broad: "Taking an action in response". Within the specification of the patent itself, In column 11, lines 32-43, we find:

> "Countermeasures range from very passive responses, such as report dissemination to other monitors 16a-16f or administrators, to highly aggressive actions, such as severing a communication channel or the reconfiguration of logging facilities within network components (e.g., routers, firewalls, network services, audit daemons). An active response may invoke handlers that validate the integrity of network services or other assets to ensure that privileged network services have not been subverted. Monitors 16a-16f may invoke probes in an attempt to gather as much counterintelligence about the source of suspicious traffic by using features such as traceroute or finger."

Both the LAN monitor and the DIDS director were capable of making reports to other monitors or administrators, and therefore fall within the scope of claim 3. It is therefore invalid.

4. The method of claim 1, wherein the plurality of network monitors include an API for encapsulation of monitor functions and integration of third-party tools.

In my opinion, adding APIs to export and import functionality of software systems is a widespread and standard practice in software engineering and should have been obvious to anyone with any professional knowledge of software development.

Furthermore, SRI was participating in the CIDF process, an open public standards effort tasked to provide APIs for intrusion detection modules to allow them to interoperate, and which began this task before the statutory bar date for these patents, which was well known in the field at the time. SRI should not be simultaneously participating in a public process to develop APIs for intrusion detection modules, and then going back after the fact and attempting to patent the idea of adding APIs to intrusion detection modules.

Therefore claim 4 is invalid on grounds of obviousness.

5. The method of claim 1, where the network is a TCP/IP network.

DIDS was deployed only on TCP/IP networks, and repeatedly refers to TCP based tools such as rlogin, telnet, and finger. Under the LAN monitor discussion [ref DIDS] it says that the LAN monitor "is able to construct higher-level objects such as connections (logical circuits), and service requests using the TCP/IP or UDP/IP protocols." This make it clear that DIDS falls within the scope of claim 5. Therefore claim 5 is invalid.

6. The method of claim 1, wherein the network monitors are deployed at one or more of the following facilities of the enterprise network: {gateways, routers, proxy servers}.

Figure 1 of [DIDS] on page 176, which is reproduced as Figure 25 of this report above, shows a monitor deployed on the gateway at the entrance to a network. Therefore, the technique of claim 8 is explicitly disclosed by DIDS and this claim is invalid for that reason alone.

Furthermore, the deployment of intrusion detection algorithms in at least routers was known in the intrusion detection community prior to November 1997, including and especially at DARPA principal investigator meetings where SRI staff were present. For example, the IDIP system [IDIP] had modules specifically designed to operate in routers and have the routers collaborate to trace the source of problems and block connections from systems that had been identified as intrusive. The JiNao system had modules intended to work inside routers [JiNao]. Chen and Levitt described a protocol for doing intrusion detection in routers to detect misbehaving ones. That paper was published in September 1997 [Che97]. SRI themselves disclosed the possibility prior to the statutory bar. For example, in the 1997 NISSC paper they write on page 355:

> "Service monitors are dynamically deployed within a domain to provide localized real-time analysis of infrastructure (e.g., routers or gateways)..."

Finally, claims 12 to 17 are materially identical to claims 1 to 6 except that they are "system" claims rather than "method" claims. It appears to me that this has no influence on the fact that DIDS fits within all of them (with obvious extensions). For the sake of brevity and clarity, I will not repeat the arguments made above for claims 1 to 6, but the reader may apply them to claims 12-17 by simply incrementing the numbers above by 11 if necessary.

The 615 patent (Patent No. 6,711,615) – DIDS invalidates.

The 615 patent claims asserted against ISS (1-6 and 13-18) do not differ materially from the 203 patent claims asserted. After a short discussion of the differences in Claim 1, I simply cross-reference back to the relevant analysis of the 203 claims, rather than repeat an identical discussion.

> *1. A computer-automated method of hierarchical event monitoring and analysis within an enterprise network comprising:*
>
> *deploying a plurality of network monitors in the enterprise network;*
>
> *detecting, by the network monitors, suspicious network activity based on analysis of network traffic data selected from one or more of the following categories: {network packet data transfer commands, network packet data transfer errors, network packet data volume, network connection requests, network connection denials, error codes included in a network packet, network connection acknowledgements, and network packets indicative of well-known network-service protocols};*
>
> *generating, by the monitors, reports of said suspicious activity; and*
>
> *automatically receiving and integrating the reports of suspicious activity, by one or more hierarchical monitors.*

The only material difference from claim 1 of the 203 patent is that the list of possible categories of network traffic has been made longer by the addition of "network connection acknowledgements", and "network packets indicative of well-known network-service protocols". Since DIDS fits entirely in the shorter list, a-fortiori it fits in the longer list.

> *2. The method of claim 1, wherein integrating comprises correlating intrusion reports reflecting underlying commonalities.*

See the analysis of the 203 patent claim 2.

> *3. The method of claim 1, wherein integrating further comprises invoking countermeasures to a suspected attack.*

See the analysis of the 203 patent claim 3.

> *4. The method of claim 1, wherein the plurality of network monitors include an API for encapsulation of monitor functions and integration of third-party tools.*

See the analysis of the 203 patent claim 4.

> *5. The method of claim 1, wherein the enterprise network is a TCP/IP network.*

See the analysis of the 203 patent claim 5.

> *6. The method of claim 1, wherein the network monitors are deployed at one or more of the following facilities of the enterprise network: {gateways, routers, proxy servers}.*

See the analysis of the 203 patent claim 6.

> *13. An enterprise network monitoring system comprising:*
> *a plurality of network monitors deployed within an enterprise network, said plurality of network monitors detecting suspicious network activity based on analysis of network traffic data selected from one or more of the following categories: {network packet data transfer commands, network packet data transfer errors, network packet data volume, network connection requests, network connection denials, error codes included in a network packet, network connection acknowledgements, and network packets indicative of well-known network-service protocols};*
> *said network monitors generating reports of said suspicious activity; and*
> *one or more hierarchical monitors in the enterprise network, the hierarchical monitors adapted to automatically receive and integrate the reports of suspicious activity.*

This is not materially different than claim 1 of this 615 patent. See the analysis there.

> *14. The system of claim 13, wherein the integration comprises correlating intrusion reports reflecting underlying commonalities.*

See the analysis of the 203 patent claim 2.

> *15. The system of claim 13, wherein the integration further comprises invoking countermeasures to a suspected attack.*

See the analysis of the 203 patent claim 3.

> *16. The system of claim 13, wherein the plurality of network monitors include an application programming interface (API) for encapsulation of monitor functions and integration of third-party tools.*

See the analysis of the 203 patent claim 4.

> *17. The system of claim 13, wherein the enterprise network is a TCP/IP network.*

See the analysis of the 203 patent claim 5.

> *18. The system of claim 13, wherein the network monitors are deployed at one or more of the following facilities of the enterprise network: {gateways, routers, proxy servers}.*

See the analysis of the 203 patent claim 6. DIDS invalidates all asserted claims of the 615 patent.

The 212 patent (Patent No. 6,708,212) – DIDS invalidates

Using SRI's claim construction and Symantec's claim construction, the 212 patent is invalidated by DIDS. Unless otherwise specified all references to DIDS in the following discussion refer to [DIDS].

> *1. Method for monitoring an enterprise network, said method comprising the steps of:*
>
> *deploying a plurality of network monitors in the enterprise network;*
>
> *detecting, by the network monitors, suspicious network activity based on analysis of network traffic data, wherein at least one of the network monitors utilizes a statistical detection method;*
>
> *generating, by the monitors, reports of said suspicious activity; and automatically receiving and integrating the reports of suspicious*

activity, by one or more hierarchical monitors.

According to SRI's claim construction, a network monitor is a

> "Process or component in a network than can analyze data; depending on the context in specific claims, the network traffic monitor may analyze network traffic data, reports of suspicious network activity, or both."

A hierarchical monitor is a

> "Process or component that receives reports from at least one lower-level monitor."

With these interpretations of the terms, it is clear that DIDS has all of the elements of a system covered by claim 1. DIDS had a LAN monitor (the NSM) which could "analyze network traffic data", and a plurality of those could be deployed in the enterprise network. On page 168 of [DIDS] it says that, "DIDS components included the DIDS director, a single host monitor per host, and a single LAN monitor for **each** LAN segment of the monitored network." (Emphasis added). This makes clear that a plurality of monitors was specifically contemplated in DIDS.

Clearly, the LAN monitor was intended to "detect suspicious network activity based on analysis of network traffic data": that was its entire purpose. Futhermore, it utilized a statistical detection method. This is specifically indicated in the DIDS reference [DIDS p 171]:

> "The LAN monitor also uses and maintains **profiles of expected network behavior**. The profiles consist of expected data paths (e.g. which systems are expected to establish communication paths to which other systems, and by which service) and service profiles (e.g. what a typical telnet, mail, or finger was expected to look like)."

And

> "The LAN monitor also uses heuristics in an attempt to identify the likelihood that a particular connection represents intrusive behavior. These heuristics consider the capabilities of each of the network services, the level of authentication required for each of the services, the security level for each machine on the network, and signatures of past attacks. **The abnormality of a connection is based on the probability of that particular connection occurring and the behavior of the connection itself.**

Thus the DIDS reference clearly discloses the use of statistical profiles to determine that network activity might be evidence of intrusion – it talks about maintaining profiles and using probabilities to decide the abnormality of particular connections.

We also know from the NSM references discussed above that in fact NSM was building long-term profiles ("network masks") using exponential averaging of the likelihood of a connection on a given day and keeping distributions for measures such as the number of bytes and number of packets in those connections, and was keeping track of a current traffic matrix and comparing it to that long-term profile.

We also know that the LAN monitors in DIDS generated DIDS events, which constitutes "generating reports of suspicious activity" and utilized the LAN Agent subcomponent to send these to the DIDS director. The DIDS director is a monitor according to the SRI claim construction since it is a "process or component that can analyze data", and specifically it analyzes "reports of suspicious network activity" – those sent to it by the LAN Monitors (and also Host Monitors). Further, the DIDS director is a hierarchical monitor, since it "receives reports from at least one lower-level monitor", to whit the LAN Monitors.

Thus DIDS is engaged in "automatically receiving and integrating the reports of suspicious activity, by one or more hierarchical monitors" since the Director was certainly integrating the reports as it ran its rules on them, decided which activities were performed by the same user as which other activities on other computers, and reported suspected intrusions to the users of the system. Hence all elements of claim 1 of the 212 patent were present in DIDS publications seven years prior to the filing of the earliest of the patents in suit. Thus, under the SRI claim construction, this claim impermissibly covers prior art and is invalid.

We now turn to whether the system of claim 1 as narrowed by claim 2 would be novel over DIDS. Claim 2 reads:

> *2. The method of claim 1, wherein at least one of the network monitors utilizes a signature matching detection method.*

DIDS was intended to be capable of detecting doorknob-rattling types of attacks. For example, on page 168 of [DIDS]

> "In a doorknob attack the intruder's goal is to discover, and gain access to, insufficiently-protected hosts on a system. The intruder generally tries a few common account and password combinations on each of a number of computers."

And

> "Because DIDS aggregates and correlates data from multiple hosts and the network, it is in a position to recognize the doorknob attack by detecting the pattern of repeated failed logins even though there may be too few on a single host to alert that host's monitor."

According to SRI's patent specification, the above detection of repeated failed logins is signature analysis being performed by the DIDS director (which is "at least one of the network monitors"). Specifically of signature analysis, in the 212 patent on column 7 line 55, it reads:

> "**Threshold analysis is a** rudimentary, inexpensive **signature analysis technique** that records the occurrence of a of specific events and, as the name implies, detects when the number of occurrences of that event surpasses a reasonable count. For example, monitors can encode thresholds to monitor activity such as the number of fingers, pings, or **failed login requests to accounts such as guest,** demo, visitor, anonymous FTP, or employees who have departed the company."

[Emphasis added.] So by SRIs own example in the specification, what the DIDS director was doing was "a signature matching detection method." This invalidates claim 2.

Additionally, in a different DIDS reference [SNA 91] there is an entire chapter of this Master's thesis that was devoted to the discussion of signature detection methods in DIDS. Detailed definitions of specific signatures for a number of attack scenarios are laid out. Nor was this only being done at the level of the DIDS Director; we know that the NSM itself detected doorknob rattling through essentially identical techniques. [Ref NSM] says that

> "The prototype is currently looking for very simple patterns: a single host communicating with more than fifteen other hosts, logins (or attempted logins) from one host to fifteen or more other hosts, and an [*sic*] any attempt to communicate with a non-existent host."

For these and other reasons, we see that the system of the 212 patent claim 1 as narrowed by claim 2 fails to distinguish any inventions from the prior art and therefore claim 2 is invalid.

> *3. The method of claim 2, wherein the monitor utilizing a signature matching detection method also utilizes a statistical detection method.*

DIDS discloses a signature detection method in the LAN Monitor at p. 12:

> "In DIDS, the monitoring and analysis functions are distributed among several components. These components include a DIDS director, a collection of host monitors, and at least one LAN monitor. The host and LAN monitors are primarily responsible for detecting single events and known attack signatures which have a high probability of being relevant to the security of a system; so they must constantly monitor their respective domains." (12)

In addition, the LAN monitor employs statistical detection methods (disclosed in [DIDS] as discussed under claim 1. Therefore, DIDS anticipates claim 3 of this patent. Furthermore, a number of references in the prior art involve combining both statistical and signature based detection in a single monitor – IDES, NIDES [NIDES], and JiNao [JiNao] are examples. It was widely understood in the community that this was possible.

> *4. The method of claim 1, wherein integrating comprises correlating intrusion reports reflecting underlying commonalities.*

The patent specification itself discloses almost nothing about what is meant by "correlating intrusion reports reflecting underlying commonalities". The only relevant text I am able to find is in column 4 line 40:

> "The enterprise monitor 16f (or monitors, as it would be important to avoid centralizing any analysis) focuses on network-wide threats such as Internet worm-like attacks, attacks repeated against common network services across domains, or coordinated attacks from multiple domains against a single domain. As an

> enterprise monitor 16f recognized commonalities in intrusion reports across domains (eg. The spreading of a worm or a mail system attack repeated throughout the enterprise 10), the monitor.." (discussion on possible responses follows)

It does not appear to me that this text enables anyone to build a system capable of "correlating intrusion reports reflecting underlying commonalities". It announces that such functionality in Emerald would be desirable, but gives no specifics on how it might be accomplished.

So, given that this is one of SRI's techniques of "correlating intrusion reports reflecting underlying commonalities", it seems to me clear that DIDS would have to fall under the same claim, because as I noted in my analysis of claim 1, on page 168 of [DIDS] we find

> "In a doorknob attack the intruder's goal is to discover, and gain access to, insufficiently-protected hosts on a system. The intruder generally tries a few common account and password combinations on each of a number of computers."

And

> "Because DIDS aggregates and correlates data from multiple hosts and the network, it is in a position to recognize the doorknob attack by detecting the pattern of repeated failed logins even though there may be too few on a single host to alert that host's monitor."

Therefore DIDS anticipates claim 4, which is thus invalid.

> *5. The method of claim 1, wherein integrating further comprises invoking countermeasures to a suspected attack.*

Since SRI includes reporting of results to people in its definition of countermeasures in the patent specification, DIDS also comes under claim 5. Specifically, in column 12 line 16, we have

> "Countermeasures range from very passive responses, such as report dissemination to other monitors 16a-16f **or administrators**, to highly aggressive actions…"

Since notifying administrators of security alerts is included in SRI's definition of countermeasures, the DIDS user interface falls under claim 5.

Even if the definition of countermeasure was restricted to aggressive measures, adding active responses to intrusion detection was under wide public discussion before November 1997, and thus doing the same in EMERALD should not be patentable. The reader is referred to that earlier discussion rather than repeat it identically here. For at least these reasons claim 5 is invalid.

> *6. The method of claim 1, wherein the plurality of network monitors includes an API for encapsulation of monitor functions and integration of third-party tools.*

Adding APIs to software in general and intrusion detection systems in particular was a commonplace obvious idea by the time of these patents. SRI was participating in open-

standards efforts to develop common APIs for intrusion detection efforts before the statutory bar data for these patents. See claim 4 of the 203 patent for a more detailed discussion of this point.

7. The method of claim 1, wherein the enterprise-network is a TCP/IP network.

DIDS was deployed only on TCP/IP networks, and repeatedly refers to TCP based tools such as rlogin, telnet, and finger. Under the LAN monitor discussion [DIDS] it says that the LAN monitor "is able to construct higher-level objects such as connections (logical circuits), and service requests using the TCP/IP or UDP/IP protocols." This make it clear that DIDS falls within the scope of claim 7. Therefore claim 7 is invalid.

8. The method of claim 1, wherein the network monitors are deployed at one or more of the following facilities of the enterprise network: {gateways, routers, proxy servers}.

Figure 1 of [DIDS] on page 176, which is reproduced as Figure 25 of this report above, shows a monitor deployed on the gateway at the entrance to a network. Therefore, the technique of claim 8 is explicitly disclosed by DIDS.

Furthermore, I argued in detail in the discussion of Claim 6 of the 203 patent that deploying intrusion detection monitors in routing infrastructure was a widely understood possibility before the statutory bar, and was thus obvious.

Therefore, claim 8 is invalid.

9. The method of claim 1, wherein deploying the network monitors includes placing a plurality of service monitors among multiple domains of the enterprise network.

10. The method of claim 9, wherein receiving and integrating is performed by a domain monitor with respect to a plurality of service monitors within the domain monitor's associated network domain

11. The method of claim 1, wherein deploying the network monitors includes deploying a plurality of domain monitors within the enterprise network, each domain monitor being associated with a corresponding domain of the enterprise network.

12. The method of claim 11, wherein receiving and integrating is performed by an enterprise monitor with respect to a plurality of domain monitors within the enterprise network.

I am dealing with claims 9-12 as a unit since they are all essentially associated with the fact that the patent specification specified a three level hierarchy of monitors in which service monitors fed domain monitors which were responsible for individual administrative domains of the network, and those in turn fed enterprise monitors. While

earlier claims did not restrict themselves to this structure, these claims do restrict themselves to it.

SRI did not invent the idea of three level hierarchies for intrusion detection systems in these patents. On the contrary, the idea of deeper hierarchies was disclosed at least in the UC Davis ISM paper [Heb97], and the GrIDS references [GrIDS & GrIDS96]. Additionally, SRI themselves disclosed it before the statutory bar both in the NIDES final report [ref] and the EMERALD NISSC paper in October 1997 [EMERALD].

The DIDS reference itself [DIDS p 174] announces an intent that "In support of the ongoing development of DIDS we are planning to extend our model to a hierarchical Wide Area Network". Then in 1992, UC Davis published a paper on a model, The Internetwork Security Model (ISM), for doing intrusion detection in large-scale networks [Heb97]. That paper discloses:

> "The ISM extends research and development efforts already existing in the field of intrusion detection. Primarily, the ISM extends the Distributed Intrusion Detection System (DIDS) [DIDS] into arbitrarily wide networks. Multiple DIDS-like monitors, called ISM domain monitors, communicating through well-defined protocols form the core of the distributed ISM."

And

> "The ISM model also allows ISMs to be grouped hierarchically. For example, ISM1 may monitor a domain which is divided into three sub-domains, each with its own ISM sub-monitors. This hierarchical structure provides two major benefits. First, because the ISM1 domain can look into its subdomains, it can aggregate a user's activities across these subdomains. This functionality is provided by additional requests which can only be made by a direct parent ISM...."

UC Davis went on to implement GrIDS, which explicitly had a hierarchy as discussed at length above, and disclosed the complete detailed design for doing so.

All of this occurred before the statutory bar on the parents in suit, and indeed before SRI began developing Emerald. Therefore, the idea that one could extend an IDS such as DIDS into a hierarchy was obvious by 1997. The fact that SRI wanted to do exactly the same to a NIDES-like system was not a novel invention, and therefore claims 9-12 are invalid.

> *13. The method of claim 11, wherein the plurality of domain monitors within the enterprise network establish peer-to-peer relationships with one another.*

There are significant difficulties with unconstrained peer-to-peer communication between intrusion detection systems because there is a risk of feedback loops in which one peer tells another peer something, and later hears back a possibly amplified version of the same thing, via some circuit of peers, causing it to be on higher alert based on its own information. Depending on exactly what kind of data is being shared, runaway chain reactions are possible without careful design.

However, I believe that this claim is invalid on the grounds that the patent does not enable a skilled person to manage peer-to-peer relationships, and certainly does not disclose anything over and above what is disclosed in the prior art, including the 1997 NISSC paper. The only reference to peer-to-peer relationships in the patent specification is on column 3, line 32:

> "Where mutual trust among domains 12a-12c exists, domain monitors 16d-16e may establish peer relationships with one another. Peer-to-peer subscription allows domain monitors 16d-16e to share analysis reports produced in other domains 12a-12c. Domain monitors 16d-16e may use such reports to dynamically sensitize their local service monitors 16a-16c to malicious activity found to be occurring outside a domain 12a-12c."

This provides no detail that would enable me to determine how the patent would solve the feedback problems mentioned above. Furthermore, an inspection of the EMERALD code provided as an appendix to the patent reveals that there is no "resource object" configuration information. Therefore, no worked examples of how a peer-to-peer relationship might actually work in Emerald are provided.

In the 1997 NISSC paper, before the statutory bar requirement, peer-to-peer relationships between domain monitors were described thusly (on page 363):

> "Where mutual trust among domains exists, domain monitors may establish peer relationships with one another. Peer-to-peer subscription allows domain monitors to share intrusion summaries from events that have occurred in other domains. Domain monitors may use such reports to dynamically sensitize their local service monitors to malicious activity found to be occurring outside the domain's visibility. Domain monitors may also operate within an enterprise hierarchy, where they disseminate intrusion reports to enterprise monitors for global correlation. Where trust exists between domains, peer-to-peer subscription provides a useful technique for keeping domains sensitized to malicious activity occurring outside their view."

The textual similarities between the language in the 1997 paper and the patent specification should be obvious, and it is also clear that the specification contains no additional detail. It appears to me that the EMERALD inventors at the time of patent filing had a general feeling that peer-to-peer relationships might be desirable, but had not actually developed any techniques for managing them usefully. There is no evidence of implementation of the techniques. If they had developed or implemented peer-to-peer notifications in EMERALD, they certainly didn't disclose them in the patent specification or appendix.

The ISM paper [ISM] provides more detail. In a section titled "ISM Peer-Level Communication", the paper says:

> "An ISM is responsible for a specific set of hosts. When a user initiates a connection from a host in one domain to a host in a second ISM domain, the ISM's may exchange information to allow a more accurate analysis for the security state of their own domains. At a minimum, an ISM must be able to identify the source (local or external to the domain) for connections leaving its

domain. If the user initiating the connection originated inside the ISM domain, the ISM need only respond that the connection began internally and not reveal the actual origin of the user. If the connection originated outside the ISM domain (e.g., the user merely passed through the domain), the ISM must respond with the host-to-host connection defintion of the connection entering the domain. This minimum capability of an ISM prevents an intruder from exploiting the domain in an attempt to disguise his origin. The protocol to support this functionality is presented below:

- "GET TIME <time>
- "GET CONNECTION TCP/IP-DEF <def> TIME <time>
- "GET ORIGIN CONN-ID <id>"

Here a specific peer-to-peer extension to DIDS is disclosed with considerably more enabling detail of mechanism than the patent specification has (the ISM paper is proposing specific message content for the peer-to-peer exchange protocol, which is lacking in the patent specification). There is motivation to combine ___ reference since the ISM paper indicates it is a proposed extension to DIDS.

Additionally, other references in the prior art disclose the idea of peer-to-peer relationships between intrusion detection systems. For example, in [ref Cooperating Security Managers], White and Pooch write:

> "The basic idea behind CSM is to have each host on a network (or internetwork) run a copy of CSM as a background process. The cooperative methodology which CSM is based on then consists of several aspects. First, CSM takes a proactive approach to both user tracking and distributed intrusion detection. Each CSM is not only concerned with what is happening to the host it is protecting, it is also concerned with actions a current user of the host performs on other systems further along the chain of connected systems. To accomplish this, CSM records information passed back along the chain on what a user is doing at other sites. Local and distributed IDSs are designed to take advantage of the information that is obtained in this manner to to detect anomalous activities. The result of this passing of information is the assurance that one system (the first system the user connects to) will have a record of all activity performed by the user. Having individual systems be responsible for collecting the information on their users means that the distributed intrusion detection activity is not centered in a single host for all users, but rather is distributed among many different systems, each host being responsible for its own users."

In short, claim 13 is yet another attempt by SRI to take ideas that were under wide discussion in the intrusion-detection community at the time, make a quick vague reference to them in the specification, and then file claims indicating it was part of their invention. The claim is invalid.

14. An enterprise network monitoring system comprising:
a plurality of network monitors deployed within an enterprise network,
said plurality of network monitors detecting suspicious network

activity based on analysis of network traffic data, wherein at least one of the network monitors utilizes a statistical detection method;

said network monitors generating reports of said suspicious activity; and

one or more hierarchical monitors in the enterprise network, the hierarchical monitors adapted to automatically receive and integrate the reports of suspicious activity.

This independent claim has no material differences from claim 1 of the 212 patent. See the discussion there. The same will go for the dependent claims that follow, for each of which I simply state the matching claim amongst the earlier claims in the 212 patent.

15. The system of claim 14, wherein the integration comprises correlating intrusion reports reflecting underlying commonalities.

See claim 4.

16. The system of claim 14, wherein the integration further comprises invoking countermeasures to a suspected attack.

See claim 5.

17. The system of claim 14, wherein the plurality of network monitors include an application programming interface (API) for encapsulation of monitor functions and integration of third-party tools.

See claim 6.

18. The system of claim 14, wherein the enterprise network is a TCP/IP network.

See claim 7.

19. The system of claim 14, wherein the network monitors are deployed at one or more of the following facilities of the enterprise network: {gateways, routers, proxy servers}.

See claim 8.

20. The system of claim 14, wherein the plurality of network monitors includes a plurality of service monitors among multiple domains of the enterprise network.

See claims 9-12.

21. The system of claim 20, wherein a domain monitor associated with the plurality of service monitors within the domain monitor's associated network domain is adapted to automatically receive and integrate the reports of suspicious activity.

See claims 9-12.

22. The system of claim 14, wherein the plurality of network monitors include a plurality of domain monitors within the enterprise network, each domain monitor being associated with a corresponding domain of the enterprise network.

See claims 9-12.

23. The system of claim 22, wherein an enterprise monitor associated with a plurality of domain monitors is adapted to automatically receive and integrate the reports of suspicious activity.

See claims 9-12.

24. The system of claim 22, wherein the plurality of domain monitors within the enterprise network interface as a plurality of peer-to-peer relationships with one another.

See claim 13.

Response to SRI's contention on DIDS and the 212 patent.

In their response to ISS's second interrogatories [Interr p18-19], SRI disputes that DIDS invalidates the 212 patent independent claims, and cite two reasons. The first is that DIDS does not disclose "use in an enterprise network, a plurality of network monitors, and a hierarchical monitor that automatically receives and integrates reports of suspicious activity from a plurality of network monitors".

The primary DIDS reference analyzed here [DIDS] does in fact disclose the use of multiple network monitors (the LAN monitor) feeding the hierarchical monitor (the DIDS director). The reader should consult the discussion of 212 claim 1 above, but the key point is that on page 168, [DIDS] discloses "a prototype Distributed Intrusion Detection System (DIDS) which generalizes the target environment in order to monitor multiple hosts connected via a network as well as the network itself. The DIDS components include the DIDS director, a single host monitor per host, and a single LAN monitor **for each LAN segment of the monitored network.**" (emphasis added). This clearly discloses the idea that the director could take input from a plurality of LAN monitors: SRI did not invent that. The paper goes on to say of the DIDS directory "thus providing the capability to aggregate information from different sources", making it clear that the DIDS director "automatically receives and integrates reports of suspicious activity from a plurality of network monitors".

SRI's second response is that DIDS does not disclose "a network monitor that utilizes a statistical detection method". On the contrary, on p 171, DIDS discloses that the LAN monitor "uses and maintains profiles of expected network behavior. The profiles consist of expected data paths (e.g. which systems are expected to establish communication paths

to which other systems, and by which service) and service profiles (e.g. what a typical telnet, mail, or finger is expected to look like)." This clearly falls within the boundaries of "statistical detection method". See the analysis of claim 1 of the 212 patent above for further detail.

Next, SRI states without any supporting argument that DIDS does not disclose the restrictions implied in claim 3, 4, 5, 6, 15, 16, and 17. The reader is referred to the discussion above of each of those claims for the reasons why they are in fact invalidated by DIDS alone or in obvious combinations with other systems.

Finally, SRI argue that

> "Defendants have also not shown that either DIDS article discloses and/or enables a plurality of service monitors, a domain monitor, or an enterprise monitor as required by claims 9-12 and 20-23."

"Service monitor", "domain monitor", and "enterprise monitor" are SRI's terms for the levels of hierarchy in a three-level hierarchical system. In our analysis of claims 9-13 above, to which the reader is referred, we pointed out that there is specific discussion in the DIDS and ISM references that suggests combining them, the combination being particularly obvious since the same research group wrote both papers, and that this obvious combination results in a multi-level hierarchy of the kind SRI is referring to. By the time of the statutory bar, UC Davis had also established with a detailed design and implementation of the GrIDS system that it was possible to take an intrusion detection system and extend it hierarchically.

IIm) Conclusions

What has been very striking to me about these patents from the first time I read them is the utter lack of respect the claims display for the prior art, even prior art of which I knew at first hand the SRI inventors were aware of. I have given presentations about GrIDS in which they sat, and sat through presentations about JiNao in which they also sat. They reference DIDS, NSM, and GrIDS in their published papers. I have worked with Phillip Porras on developing common public standards for intrusion detection systems to be modularly integrated in a building-block manner. The papers on these prior-art systems are widely cited in the research literature. And yet the claims of the patents display no effort that I can discern to work around that prior art.

III) Data and Other Information Considered in Forming the Above Opinions

IIIa) General References

[Amo94] Edward Amoroso, *Fundamentals of Computer Security Technology*, Prentice Hall. 1994

[Bac00] Rebecca Bace, *Intrusion Detection,* Macmillan, 2000.

[DOD85] *Department of Defense Trusted Computer System Evaluation Criteria*, DoD Standard DOD 5200.280STD, December, 1985.

[Kar02] P. Karger and R. Schell. *Thirty Years Later: Lessons from the Multics Security Evaluation*, Proceedings of the 2002 Annual Computer Security Applications Conference.

[Rit79] D. Ritchie, *The Evolution of the Unix Time-Sharing System*. Proceedings of the Language Design and Programming Methodology conference at Sydney, Australia, September 1979.

IIIb) Papers on Intrusion Detection

[And80] James Anderson, *Computer Security Threat Monitoring and Surveillance*, Report on Contract 79F296400, James P. Anderson Co.

[And94] D. Anderson et al., *Next Generation Intrusion Detection Expert System (NIDES), Software Users Manual – Beta-Update Release,* SRI Technical Report, December 1st, 1994.

[And95] D. Anderson et al., *Next-generation Intrusion Detection Expert System (NIDES) A Summary*, Technical Report SRI-CSL-95-07, May 1995, SRI International.

[Che97] S. Cheung and K. Levitt, *Protecting Routing Infrastructures from Denial of Service Using Cooperative Intrusion Detection,* Proc. New Security Paradigms Workshop 1997, Cumbria, UK, September 23-26, 1997. Available from http://seclab.cs.ucdavis.edu/papers/nsp.pdf.

[CIDF] S. Staniford-Chen, et al., *The Common Intrusion Detection Framework - Data Format,* available at http://dougmoran.com/tatzlwyrm/CACHE/cidf-rfc.txt.

[Comer] Douglas E. Comer, *Interworking with TCP/IP,* Vol. 1 (5th Edition) (June 30, 2005).

[Den86] Dorothy Denning, *An Intrusion Detection Model*, Proceedings of the 1986 IEEE Symposium on Security and Privacy, Oakland, 1986.

[DIDS] S.R. Snapp, J. Brentano, G.V. Dias, L.T. Heberlein, C. Ho, K.N. Levitt, B. Mukherjee (with S.E. Smaha, T. Grance, D.M. Teal, D.L. Mansur), "*DIDS - - Motivation, Architecture, and an Early Prototype,*" Proc. 14th National Computer Security Conference, Washington, DC, Oct. 1991, pp. 161-176.

[DIDS-Feb] Steven R. Snapp et al., "*Intrusion Detection Systems (IDS): A Survey of Existing Systems and A Proposed Distributed IDS Architecture,*" (February 1991).

[DRA] Calvin Ko, et al., *Analysis of an Algorithm for Distributed Recognition and Accountability Conference on Computer and Communications Security, Proceedings of the 1st ACM Conference on Computer and Communications Security*, Fairfax, VA 1993.

[EMERALD] P. Porras and P. Newman, (*EMERALD: Event Monitoring Enabling Responses to Anomalous Live Disturbances*), 20th NISCC, October 9, 1997.

[GrIDS] Steven Cheung, Rick Crawford, Mark Dilger, Jeremy Prank, Jim Hoagland, Karl Levitt, Stuart Staniford-Chen, Raymond Yip, Dan Zenkle, "(*The Design of GRIDS: A Graph-Based Intrusion Detection System*)," Technical Report, UC Davis Department of Computer Science, Davis California (May 14, 1997).

[GrIDS96] Chen, S.S., et al., "GrIDS-A Graph Based Intrusion Detection System For Large Networks", 19th National Information Systems Security Conference, 1996.

[Heb90] L. Todd Heberlein, et al., *A Network Security Monitor*, Proceedings of the 1990 IEEE Symposium on Security and Privacy, Oakland, 1990.

[Heb91a] L. Todd Heberlein, (*Towards Detecting Intrusions in a Networked Environment*), UC Davis Division of Computer Science Report #CSE-91-23. June 1991.

[Heb91b] L. T. Heberlein, B. Mukherjee, K. N. Levitt, *A Method to Detect Intrusive Activity in a Networked Environment*, Proc. of the 14th National Computer Security Conference, October 1991, pp. 362-371.

[Heb91c] L.T. Heberlein et al., *Towards Detecting Intrusions in a Networked Environment.* Proc. of the 14th Department of Energy Computer Security Group Conference, May 1991, pp.(17)47-(17)65.

[Heb95] L. Todd Heberlein, *Network Security Monitor: Final Report*, Lawrence Livermore National Laboratory project deliverable from UC Davis.

[Heb97] L. Todd Heberlein, et al., *Internetwork Security Monitor,* Prog. of the 15th National Computer Security Conference, October 1992, pp. 262-271.

[IDIP] Boeing Defense & Space Group, (*Intruder Detection and Isolation Protocol Concept*), Dynamic Cooperating Boundary Controllers Programs, January 1997 (ISS 27577-ISS27627).

[Interr] SRI International, Inc.'s Responses To Defendant ISS-GA's Second Set Of Interrogatories [Nos. 19-20] and SRI's Third Supplemental Response To ISS-GA's Interrogatory No. 17.

[Jav91] H. Javitz and A. Valdes, *The SRI IDES Statistical Anomaly Detector*, Proceedings of the 1991 IEEE Symposium on Security and Privacy.

[Jav93] H. Javitz and A. Valdes, *The NIDES Statistical Component: Description and Justification*, SRI Technical Report. March, 1993.

[JiNAO] Y. Frank Jou et al., (*Architecture Design of a Scalable Intrusion Detection System for the Emerging Network*), Technical Report CDRL A005, DARPA Order No. E296, Dept. of Computer Science North Carolina State University, April 1997.

[Joint Claim Construction Statement] Joint Claim Construction Statement as submitted March 17, 2006.

[Muk94] B. Mukherjee, L. T. Heberlein, K. Levitt. *Network Intrusion Detection.* IEEE Network, May-June 1994. Vol. 8 No. 3 pp 26-41.

[NetRanger] NetRanger User's Guide 1.3.1, WheelGroup Corporation, 1997.

[Por97] P. Porras and P. Neumann. *EMERALD: Event Monitoring Enabling Responses to Anomalous Live Disturbances.* Proceedings of the 20th NIST-NCSC National Information Systems Security Conference. Baltimore, October 1997.

[RealSecure 1.0] RealSecure 1.0: User Guide and Reference Manual (1996).

[RealSecure 1.1] RealSecure 1.1: User Guide and Reference Manual (1997).

[RealSecure 1.2] RealSecure 1.2: User Guide and Reference Manual (1997).

[RealSecure 1.2.2] RealSecure 1.2.2: User Guide and Reference Manual (Sept. 1997).

[RFC 2328] RFC 2328-OSPF Version 2 available at http://www.faqs.org/rfcs/rfc2328.html.

[RFC 854] RFC 854 -Telnet Protocol Specification available at http://www.faqs.org/rfcs/rfc854.html.

[RTARR] (*Real-Time Attack Recognition and Response: A Solution For Tightening Network Security*) (ISS 357242-357259).

[Santa Cruz PI] GrIDS, Graph-Based Intrusion Detection Presentation, August 27, 1996 (ISS_02126250-ISS_02126265).

[Sna 91] Steve Snapp, "Signature Analysis and Communication Issues in a Distributed Detection System," M.S. Thesis, Division of Computer Science, University of California, Davis, August 1991 (ISS 03336-03378).

[SRI06] SRI Inc., *History*, Document on history of intrusion detection research at SRI as accessed from URL http://www.csl.sri.com/programs/intrusion/history.html on 3/22/06.

IIIc) Deposition Transcripts.

[Jou06] Deposition of Frank Jou, January 27th, 2006. Raleigh, North Carolina.

IV) Trial Exhibits

I may rely on visual aids and demonstrative exhibits that demonstrate the basis of my opinions.

V) Compensation

I am being compensated for the work done in connection with this matter at a rate of $300 per hour for general consulting and report writing, and at a rate of $400 per hour for deposition and trial appearances, and preparation therefore.

VI) Previous Testimony

I have not testified as an expert at trial or by deposition within the last four years.

VII) Reservations of Rights

I reserve the right to amend or supplement the statement based on further discovery and preparation in this action, including my review of any expert statement submitted on behalf of SRI, and review of material currently designated confidential.

Dated: April 21, 2006

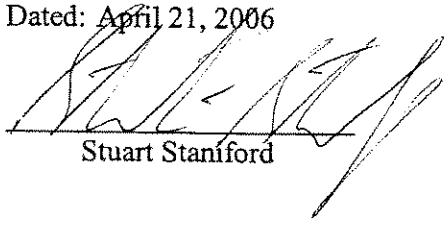

Stuart Staniford

Appendix A: Curriculum Vitae of Stuart Staniford, PhD

690 Hearst Ave,
San Francisco, CA 94112
stuartstaniford@sbcglobal.net
(415) 239-2090

Summary

I am a researcher, inventor, and entrepreneur. My scientific interests include computer worm propagation, defenses against computer worms, and issues surrounding energy security. My research has led to over one thousand academic citations to date, as well as media coverage in national magazines and newspapers. I also have business and management experience as an entrepreneur in the information security industry.

Work Experience

Invicta Consulting. President, Jan 2005 - Present

Solo Consultant. Clients presently include:

- ***King and Spalding/ISS***. Jan 2005 - present. Consult on technical issues related to pending intellectual property litigation (SRI is suing ISS for patent infringement).
- ***Nevis Networks***. Jul 2005 - present. Consult on maintenance issues arising in systems that I designed, help with patent filings arising from my time there as an employee, review security algorithms, help with product testing, and write white papers. *As of this writing, patent filing role has concluded.*
- ***FireEye***. Mar 2006 - present. Review algorithms, provide third party validation of effectiveness of system. *NB. As of this writing, engagement is verbally agreed but contract is still to be signed. Thus client engagement is likely but not certain.*

The Oil Drum. Editor, September 2005 – Present

Wrote hundreds of blog posts for this popular website exploring the scientific issues surrounding peak oil, economic response to oil shocks, and climate change.

Nevis Networks. Principal Scientist, April 2004 – July 2005

Architected a very high-speed event correlation system and the traffic anomaly subsystem for this startup developing 10Gbps network security solutions for the ethernet edge of internal enterprise networks. These systems resulted in four patents including a novel multi-dimensional external memory algorithm for storing log-records on disk at very high speeds. Designed portions of the

product's graphical user interface. Worked extensively with engineering teams in Pune, India and Santa Clara, California implementing my designs.

Silicon Defense. Founder and President, 1998 - 2004

Managed 23 staff performing a mixture of government contract research and commercial product development. Obtained ten research contracts for the company up to $2.3m in size, working for four different DARPA program managers. Performed and published research into intrusion detection, intrusion correlation, and especially worms. Work was covered in Business Week, Federal Computer Week, PC World, Network World, American Banker, and others. Spoke to a variety of audiences on research and risks in cyberspace. Coauthored patent application on invention of worm containment.

Wrote business plan for the company and raised angel capital. Sold commercial products into Fortune 500 accounts (company gained over 50 commercial customers during my tenure). Had profit and loss responsibility for a $1.75m operation. Extensive experience interacting with press, analyst, and investment communities. Also served on a number of program committees and led two standards groups.

After five years of 100%+ growth out of cashflow alone, company was obliged to file bankruptcy due to DARPA's decision to classify further research in the information security area.

UC Davis. Researcher, 1994 – 1997, and Assistant Adjunct Professor, 1997 - 1999

Founded and cochaired the working group that developed the Common Intrusion Detection Framework at the request of DARPA. This involved working with a team of over a hundred researchers and developers from a wide variety of companies and organizations. Led a team of ten researchers and students building a large, distributed, intrusion-detection system (GrIDS). Performed research in new statistical techniques to help in tracing intruders across the Internet. Presented work at conferences and to funding agencies. Wrote successful funding proposals and published papers on work.

Education

M.S. (Computer Science). March 1995. University of California at Davis. Advisor: Prof. Karl Levitt

Ph.D. (Physics) June 1993. University of California at Davis. Awarded fellowships for three years consecutively.

M.S. (Physics) June 1990. University of California, Davis.

B.Sc. (Mathematical Physics) June 1988. University of Sussex, UK. First Class Honors.

Refereed Publications

S. Staniford, D. Moore, N. Weaver, and V. Paxson, *The Top Speed of Flash Worms.* Proceeding of the ACM Workshop on Rapid Malcode (WORM), 2004

N. Weaver, D. Ellis, S. Staniford, and V. Paxson, *Worms vs Perimeters – The Case for Hard-LANS.* Proceedings of Hot Interconnects, 2004

N. Weaver, V. Paxson, and S. Staniford, *Very Fast Scanning Worm Containment.* Proceedings of USENIX Security, 2004

S. Staniford. *Containment of Scanning Worms in Enterprise Networks*. To appear in the Journal of Computer Security.

N. Weaver, V. Paxson, S. Staniford, and R. Cunningham *A Taxonomy of Computer Worms*. Proceedings of the ACM Workshop on Rapid Malcode (WORM). Washington D.C. October, 2003.

D. Moore, V. Paxson, S. Savage, C. Shannon, S. Staniford, and N. Weaver, *Inside the Slammer Worm*, IEEE Security and Privacy, July/August 2003.

S. Staniford, J. Hoagland and J. McAlerney. *Practical Automated Detection of Stealthy Portscans.* Journal of Computer Security. Vol 10, Issue 1/2, 2002.

S. Staniford, V. Paxson, and N. Weaver *How to Own the Internet in Your Spare Time,* Proceedings of the 11th USENIX Security Symposium 2002.

D. Donoho, A. Flesia, U. Shankar, V. Paxson, J. Coit, and S. Staniford, *Multiscale Stepping-Stone Detection: Detecting Pairs of Jittered Interactive Streams by Exploiting Maximum Tolerable Delay*, Proc. RAID 2002.

J. Hoagland, and S. Staniford, *Viewing IDS alerts: Lessons from SnortSnarf.* Proceedings of DISCEX II, Anaheim, June 2001.

J. Coit, S. Staniford, and J. McAlerney. *Towards Faster Pattern Matching for Intrusion Detection: Exceeding the Speed of Snort*. Proceedings of DISCEX II, Anaheim, June 2001.

R. Feiertag, S. Rho, L. Benzinger, S. Wu, T. Redmond, C. Zhang, K. Levitt, D. Peticolas, M. Heckman, S. Staniford-Chen, and J. McAlerney, *Intrusion Detection Inter-component Adaptive Negotiation.* Computer Networks. 2000.

S. Staniford, J. Hoagland and J. McAlerney. *Practical Automated Detection of Stealthy Portscans.* Proceedings of the ACM CCS IDS Workshop,November 1, 2000. Athens, Greece.

R. Feiertag, S. Rho, L. Benzinger, S. Wu, T. Redmond, C. Zhang, K. Levitt, D. Peticolas, M. Heckman, S. Staniford-Chen, and J. McAlerney,. *et al.* *Intrusion Detection Inter-Component Adaptive Negotiation.* Proceedings of the 2nd International Workshop on Recent Advances in Intrusion Detection (RAID 99), Lafayette, Indiana; September, 1999.

S. Staniford-Chen, B. Tung, and D. Schnackenberg,. *The Common Intrusion Detection Framework (CIDF).* Proceedings of 1998 Information Survivability Workshop – ISW'98, Orland, Florida; October, 1998.

S. Staniford-Chen, S. *et al* *GrIDS: A Graph-Based Intrusion Detection System for Large Networks.* Proceedings of the 19th NISSC, Baltimore, 1996.

S. Staniford-Chen, and L.T. Heberlein, *Holding Intruders Accountable on the Internet*. Proceedings of the 1995 IEEE Symposium on Security and Privacy, Oakland, CA. 1995.

J. Kiskis and S. Staniford-Chen, *Universal Amplitude Ratios and Functions for the SU(2), Finite-Temperature Phase Transition*. In Axen, D., Bryman, D., and Comyn, N. (eds) Vancouver Meeting. Particles and Fields '91. p 821. World Scientific. 1992.

Published Reports and Theses

N. Weaver, V. Paxson, and S. Staniford, *The Worst Case Worm*. Silicon Defense Technical Report. August 2003.

S. Staniford and C. Kahn, *Worm Containment on the Internal Network.* Silicon Defense Technical White Paper. March 2003

N. Weaver, V. Paxson, S. Staniford, and R. Cunningham, *Large Scale Malicious Code: A Research Agenda*. Silicon Defense Technical Report, Dec 2002.

B. Tung, et al. The Common Intrusion Detection Framework Specification. Nov 2001.

S. Staniford, O.S. Saydjari, and K. Williams *The US is Not Safe in a Cyberwar.* Paper presented to Department of Defense and National Security Council executives. May 2001. 2nd Edition.

S. Staniford, O.S. Saydjari, and K. Williams *The US is Not Safe in a Cyberwar.* Paper presented to DARPA. Sep 2000.

S. Cheung, S. *et al* *The Design of GrIDS: A Graph-Based Intrusion Detection System*. UCD Technical Report CSE-99-2, January, 1999.

S. Staniford-Chen, *Distributed Tracing of Intruders*. Master's Thesis, University of California at Davis. 1995.

S. Staniford-Chen, Finite Size Scaling and the Universality Class of SU(2) Lattice Gauge Theory. PhD Thesis, University of California at Davis. 1993.

S. Staniford-Chen, *Finite Size Scaling of Probability Distributions in SU(2) Lattice Gauge Theory and Phi^4 Field Theory* . Preprint UCD-92-17, University of California at Davis. 1992.

Patent Filings

S. Staniford and M. Bakshi, *A System and Method for Selecting Memory Locations for Overwrite.* Filed January 23rd, 2006

S. Staniford *et al*, *A System and Method for Aggregating and Consolidating Security Event Data.* Filed November 26th, 2005

S. Staniford and T. Mustafa, *A System and Method for Deprioritizing and Presenting Data.* Filed November 4th, 2005

S. Staniford and P. Sobel. *System and method for storing multi-dimensional network and security event data.* Filed October 14th, 2005

S. Staniford, C. Kahn, N. Weaver, C. Coit, and R. Jonkman , *Method and system for reducing the rate of infection of a communications network by a software worm.* Filed December 6th, 2002. Filing serial number: 313623.

Software Systems

CounterMalice was the first automated worm containment system in the world capable of containing zero-day worms, and became a commercial product. It operates by dividing a network into cells, recognizing wormlike behavior, and suppressing spread of a worm from one cell to another. CounterMalice was developed with Cliff Kahn, Nick Weaver, Jason Coit, Roel Jonkman, Joe McAlerney, and Dan Watson. My role was providing the initial vision, developing quantitative methods for tuning the system such that its performance against worms could be engineered in advance, and coding portions of the user interface.

Spice was the first system capable of detecting stealthy portscans from multiple sources using simulated annealing to correlate disparate events. It became part of a commercial product (CounterStealth). Spice was developed with James Hoagland and Dan Watson. My role was initial vision, much of the design, and techniques for validating its performance.

Spade was a network anomaly detection system (used as an input to Spice). It became well known and gained widespread operational use when it was incorporated as a plug-in into the open-source GPL intrusion detection system Snort. Spade was developed with James Hoagland. My role was the basic idea and much of the design.

Snortsnarf was an open-source alert viewer for Snort, that was innovative in systematically taking account of the possibility of attackers deliberately targeting the user interface screen real-estate. Snortsnarf gained widespread operational use at sites generating large volumes of Snort alerts, and was the main user interface for intrusion detection at the 2002 Winter Olympics. Snortsnarf was developed with

James Hoagland. My role was to build the first version of the system, and provide design input during ongoing maintenance and extension.

GrIDS was the first intrusion detection/correlation system capable of correlating alerts hierarchically to infer the presence of large scale automated attacks throughout a network (including scans and worms). The system could handle a wide variety of inference tasks through a set of rules that assembled activity into distributed graphs which the system reasoned about. The inference hierarchy could by dynamically rearranged via a drag-and-drop UI. GrIDS was developed with Mark Dillinger, James Hoagland, Chris Wee, Dan Zerkle, Rich Crawford, Steven Templeton, Stephen Cheung, and Karl Levitt. My role was that of team leader/group facilitator, contributor to the design of the inference mechanism and hierarchy, and implementer of the components that supported the rearrangable hierarchy. GriDS was tested in a medium-sized deployment at UC Davis.

Funding Obtained

Network Associates (subcontract under DARPA contract). *Intrusion Detection InterComponent Adaptive Negotiation.* $100k (1998-1999)

University of California, Davis (subcontract under DARPA contract). Global Guard: A Protection Architecture for Survivability of Large Scale, High-Confidence Information Networks. $90k (1999-2000)

The Boeing Company (subcontract under DARPA contract). Multi-Community Cyber Defense. $480k (1999-2002)

EMC Corp. Explorations of Randomness in Hard Disk Rotation Times. $32k (1999-2000)

WetStone Technologies (subcontract under DARPA contract). NetFlare IDWG subcontract. 2000-2001

DARPA Internet Trap-and-Trace. $2.3m (2000-2003). With Felix Wu (UC Davis) and Vern Paxson, ICIR.

US Air Force, Rome Labs. IDS Correlation Using IDWG. $50k (2000-2001)

US Air Force, Rome Labs. IA-INTER-OP IETF IDWG. $145k (2001-2003). Co-PI with Joseph Betser

Northrop Grumman (subcontract under DARPA contract). International Coalition Exercises. $203k (2002-2003)

BBN Technologies (subcontract under DARPA contract). Information Assurance Operational Experimentation. $500k (2002-2003)

Service

Program committee member of the **ACM Workshop on Rapid Malcode (WORM)**. 2005

Advisory board member for the **Collaborative Center for Internet Epidemiology and Defenses**. 2004-present

General chair and program committee member of the **ACM Workshop on Rapid Malcode (WORM)**. 2003

Co-organizer of the DIMACS Workshop on Large Scale Wttacks, 2003.

Served on the program committee of the Symposium on **Recent Advances in Intrusion Detection (RAID)** from 1999-2003.

Member of the **Mitre CVE Editorial Board**. This group developed a standard naming system for computer vulnerabilities. (1999-2002, now emeritus member)

Founded and cochaired the **IETF working group IDWG** (1999-2004). This working group has almost completed work on a set of documents to allow common reporting by disparate intrusion detection systems. These should become RFCs shortly.

Founded and chaired the **Common Intrusion Detection Framework** working group, at the request of DARPA (1998-2000). This group was responsible for developing a standard for all DARPA-funded intrusion detection researchers to build their systems to in order to allow inter-operation.

Invited Presentations

Usenix Security, 2004. *Military Strategy in Cyberspace.* San Diego, August 2004.

University of California, Davis. *Worms and Worm Containment.* Seminar at Computer Science Department, Feb 2003.

John Moores University Computer Science Department. *Worms and Worm Containment*. Seminar at Computer Science Department, Dec 2003.

DIMACS Workshop on Large Scale Attacks. *Introduction to Worms and Worm Containment.* Oct 2003.

The Forum on Information Warfare. Future *Technologies of Cyberwar Operations*. November 2003

Microsoft Corporation. *Worms and CounterMalice – presentation to the Security Business Unit*. Sep 2003.

Government Communications Conference. *Cyber-Weapons of Mass Destruction.* Invited Keynote Presentation. July 2003.

Annual Computer Security Applications Conference. *Defeating Worms.* Invited panel presentation. Dec 2002.

AT&T. *Worms and Anti-worm devices.* Invited presentation to security group. Sep 2002.

National Security Agency. *Worms and Traceback.* Invited presentation to technical groups. Aug 2002.

UC Berkeley. *Military Strategy in CyberSpace.* Invited lecture as part of a special series of lectures on critical infrastructure protection. Mar 2002.

Ground Systems Architectures Workshop. *CyberSpace risks to Ground Systems.* Invited presentation on risks to satellite ground systems due to dependence on the Internet. Mar 2002.

Annual Computer Security Applications Conference. *IDWG Progress Report* – invited panel presentation. Dec 2001.

ACM Conference on Computer Security. *Detecting Distributed Portscans.* Tutorial as part of joint tutorial with Vern Paxson on Intrusion Correlation. Nov 2001.

RAID Symposium. *State of Intrusion Detection.* Invited Panel Presentation. Oct 2001.

National Security Telecommunications Advisory Council. *The US is not safe in a cyberwar.* Joint work with O. Sami Saydjari (presenting) and Ken Williams. June 2001.

SRI Workshop on Adversary Characterization. *Cyberwar and Strategy – some lessons from history.* Aug 2001.

SANS National Conference. *Viewing Snort Alerts with Snortsnarf.* May 2001.

CanSecWest. *Spade and Spice.* Mar 2001.

RAID Symposium. *IDWG: Progress towards an open IDS alert standard.* October 2000.

National Security Council. *Presentation to members of the NSC staff on future risks from cyber attacks on US.* Sep 2000.

RAID Symposium. *IDS Standards – Lessons Learned to Date.* September 1999.

CIO Council, Monterey Meeting. *Standardizing IDS Alerts*. March 1999.

White House Workshop on Cybersecurity Research. *Standardizing IDS Alarms*. February 1999.

Selected Press Coverage of Work

Stuart Staniford's work has been featured in several dozen news stories in the major media and computer technical publications. A small sample include:

Business Week To Trap a Superworm
http://www.businessweek.com/technology/content/feb2003/tc20030225_4104_tc047.htm
The Slammer worm's ability to spread so rapidly adds a frightfully new dimension to the species. Does Stuart Staniford have the cure?

PC World Dawn of the Superworm,
http://www.pcworld.com/news/article/0,aid,110014,00.asp

ComputerWorld Study: Slammer was fastest spreading worm yet,
http://www.idg.com.hk/cw/readstory.asp?aid=20030205005

The Independent Internet worm took 10 minutes to create global chaos,
http://news.independent.co.uk/digital/news/story.jsp?story=375374

CERTIFICATE OF SERVICE

I, Ryan J. Stempniewicz, hereby certify that a true and complete copy of the Expert Report of Stuart Staniford has been caused to be delivered to SRI International, Inc. and Symantec Corporation on this 21st day of April 2006 in the manner indicated:

REPORT WITHOUT EXHIBITS VIA EMAIL;
REPORT WITH EXHIBITS BY U.S. POSTAL SERVICE

Howard G. Pollack
Fish & Richardson P.C.
500 Arguello Street, Suite 500
Redwood City, CA 94063

Paul S. Grewal
Day Casebeer Madrid & Batchelder LLP
20300 Stevens Creek Boulevard
Suite 400
Cupertino, CA 95014

REPORT WITH EXHIBITS BY U.S. POSTAL SERVICE

Timothy Devlin
Fish & Richardson P.C.
919 N. Market Street, Suite 1100
P.O. Box 1114
Wilmington, DE 19899

Richard K. Herrmann
Morris James Hitchens & Williams LLP
222 Delaware Avenue, 10th Floor
P.O. Box2306
Wilmington, DE 19899-2306

_____
Ryan J. Stempniewicz, Esq.